Transcribed police scanners in real-time (murph.live)

211 points by illuminated 1 day ago | hide | past | web | favorite | 100 comments

add comment

VWWHFSFQ 1 day ago [–]

The speech-to-text transcription is so incredibly wrong that it's almost dangerous to publish it like this.

For instance:

> at the beach view new screen for assistance there is a needle in his hand he's foaming from his mouth throwing off this item

What the officer actually said on the radio:

> He was going to Rainier Beach area. A request for assistance to approach two people with needles. <operator>: Call was from a neighbor in the area.

reply

lyap 1 day ago [–]

> I know up until injuries this was in France but when I press your boobs

> Jun 9, 2020 12:17 PM

Sounds like it's getting them very wrong.

reply

runawaybottle 1 day ago [–]

Well, when you come realize the training data for a lot of this is often someone reading the Washington Post corpus or something similar, it kind of makes sense why accuracy kind of sucks.

reply

ghostpepper 1 day ago [–]

It would be cool if there was a way to listen for a bit and feed some corrected transcriptions back in to help train the algorithm better

reply

Cthulhu_ 1 day ago [–]

Yeah, live transcription is one of those jobs that can be easily done remotely, at scale, and crowdsourced; if one sentence is transcribed by two or three people, you can do error correction / checking as well (or just show the different interpretations). Transcripts for radio comms are important because they can be used in legal proceedings. Same with e.g. bodycam footage.

reply

watwut 1 day ago [–]

You cant use crowd sourced transcript for legal proceedings if original audio was not recorded. It would be too easy to be poisoned by bad actors doing intentionally bad transcripts.

reply

hammock 23 hours ago [–]

I guess it would be considered hearsay at that point

reply

MengYuanLong 1 day ago [–]

I don't have familiarity with speech-to-text but wouldn't it be possible to weight words based on their probability in this application to help resolve this. For example, suspect is probably a low frequency word in normal speech but very high in radio chatter.

reply

amluto 1 day ago [–]

This needs care. Imagine if you accidentally trained a model that added racism when none was present in the audio.

reply

VWWHFSFQ 1 day ago [–]

Or used a weird model like this one which does some sort of Markov chaining to complete sentences that weren't even present in the original audio or transcription.

"foaming at the mouth" was never even close to being uttered on the radio. I'm guessing the (flawed) model inserted that part because of the proximity to the word "needle" and "assistance".

Maybe? No idea.. this website it totally fucked.

reply

optimumsuave 1 day ago [–]

Hello!

The quality is currently limited by Google's API. I am working on getting some pre-trained models implemented, but voice processing is not my speciality as a software engineer.

I do NOT want to spread misinformation nor do we want to unjustly slander anyone. Tonight I will be adding a disclaimer mentioning the limitations of our service and will make sure it is forefront on the website.

Hopefully we can create a model which can deliver better results.

reply

rasz 1 day ago [–]

apple dictation does this, it will change already transcribed text mid sentence if it thinks something else fits better

https://news.ycombinator.c...

at 33 second mark https://twitter.com/jamesc...

reply

n4r9 1 day ago [–]

That's slightly different to what OP was talking about, if I'm understanding correctly. You're talking about reassessing the probability of previous words based on future words. They're talking about weighting the prior probability of each word based on the context i.e. police conversation as opposed to normal phone conversation.

reply

marakv2 1 day ago [–]

[my side project is using speech-to-text so this is amateur hour response]

Yes, but it takes a bit to build up that weighted list and it can be quite hefty to parse. So likely they may be building this behind the scenes currently. As another commenter pointed out, being able to correct a chunk and send it back to help the algorithm would be a nice feature here.

Side note: I'm dealing with this issue at the moment - if anyone has a good resource on reducing the workload I'd love a link!.

Ed:spelling

reply

watwut 1 day ago [–]

Imo, text to speech that produces apparent garbage is better then one who produces probably stuff that is wrong. Someone, either cop or citizen, could easily end up accused of wrongdoing where not actual wrongdoing happened.

reply

haram_masala 1 day ago [–]

You could actually borrow some techniques from text mining to do this, e.g. probabilistic latent semantic analysis, to constantly re-train your speech recognition model by reinforcing translations that semantically "make sense."

reply

abngb 1 day ago [–]

Traditional Automatic Speech Recognition (ASR) systems do this, and this component is known as a Language Model (LM).

Typically you would use/train a LM for your domain or specifically for your dataset.

reply

foobiekr 1 day ago [–]

if it wasn't so serious, some of them would be are sort of entertaining.

reply

asclimo 1 day ago [–]

* For entertainment purposes only.

reply

blantonl 1 day ago [–]

This is very impressive.

I'm the owner of Broadcastify.com, where presumably these streams are being transcribed from. We've dabbled in this space and looked at real-world approaches to taking something like this to market, but transcribing 7000+ streams to text seems like an expensive (computational) and ($$) effort that needs a lot of investigation.

Note to mention that the individual lexicons between streams are drastically different.

I wonder how the developer has done the integration to our streams... because I never heard from them :)

reply

lunixbochs 1 day ago [–]

I prototyped this concept too, at https://feeds.talonvoice.com with prohibitively expensive Google speech recognition, but also have a feature for users to listen and fix transcriptions. If murph was anything like me they probably paid for broadcastify and tailed a couple of the static mp3 feeds.

My plan was to collect user transcription corrections on my site then train my own inexpensive models on them. The open-source speech tech I work on can do passable transcription at close to 10x faster than realtime on a quad core desktop CPU (or 200 simultaneous streams per 4-core box at 50% activity). With higher quality transcription it's closer to 10-20x faster than realtime.

For your case you could also try to move most of the computation down to the uploading machine. These models can run on a raspberry pi.

I think the biggest work for a new effort here is going to be building local language models and collecting transcribed audio to train on. However, there have been a couple of incredible advances in the last year for semi-supervised speech recognition learning, where we can probably leverage your 1 year backlog as "unsupervised training data" while only having a small portion of it properly transcribed.

The current state-of-the-art paper uses around 100 hours of transcribed audio and 60,000 hours of unlabeled audio, and I bet you could push the 100h requirement down with a good language model and mixing in existing training data from non-radio sources.

reply

jcims 1 day ago [–]

Not to be 'that guy' but I vastly prefer your implementation. Having both the audio and transcription is almost mandatory to something like this (unless I'm an idiot and missing the ability to play the call on this).

I wonder if one could mix in openstreetmap data for a location to help pick up local references. (Eventually would be cool to round trip it with a little ping when addresses/businesses are referenced).

reply

lunixbochs 1 day ago [–]

Yes, I think local language modeling would be crucial to doing this correctly

reply

blantonl 1 day ago [–]

Our new project, Broadcastify Calls, might be a better fit for this. Instead of 24x7 live streams, we capture and ingest every individual call as a compressed audio file from SDRs (software defined receivers) We can then ingest and present back to consumers playback, rewind, playlist, of those calls. We're now capturing over 100 systems and 800-900 calls a minute... as we solidify the architecture it will be our new direction for how we capture and disseminate public safety audio (Police Scanners)

https://www.broadcastify.c...

reply

jcims 1 day ago [–]

Love the idea! P25 decoder seems like it needs a little tuning...can you share what you're using?

Any thoughts on adding the ability to comment/transcribe/etc?

reply

p0sixlang 1 day ago [–]

Hey Blanton, big fan, and software engineer here. Any way you could add Rochester, NY (Monroe County Sheriff, and RPD) to the list of supported calls? I have an RTL SDR, but haven't been able to spend the time figuring out how to decrypt the Phase II trunking.

reply

blantonl 1 day ago [–]

You can get started as a calls ingest provider here:

https://wiki.radioreference....

reply

robotastic 1 day ago [–]

Hop on to https://gitter.im/trunk-reco... if you are having trouble getting the https://github.com/robotas... software running. Trunk Recorder puts a wrapper around the OP25 software and lets you capture all of the audio from a radio system using an SDR.

reply

lunixbochs 1 day ago [–]

The source repo to feeds.talonvoice.com includes a test ingestor that scrapes your calls API and uploads the src/dst info with the transcription, I haven't tested it live though.

reply

optimumsuave 1 day ago [–]

Hi lunixbochs!

Your prototype is amazing! The quality of transcription is definitely better than ours via Google.

After we did some legal research we wanted to avoid storing the recordings and rather solely transcription text. Giving access to a platform for humans to verify the transcriptions and in turn train the model is a great idea.

I have started working on getting some pre-trained models set up. I am trying to implement them with wav2letter, deepspeech, kaldi, vosk, etc. - I just need to be pointed in the right direction.

Raspberry Pi's were something I was considering as well - small energy footprint and powerful enough to run these models.

Do you have any advice on ML or acoustic models to avoid? I am working with the 100 hour dataset now.

Thanks!

reply

jshann8384 22 hours ago [–]

I have the same setup as Broadcastify Calls (trunkrecorder) and a site built to play each audio recording then allow the user to provide what they heard. I used it to train some public safety specific models on Kaldi and Sphinx.

I have 30ish streams and keep 6 days worth, I could keep longer if you'd like to work together on this. I reached out to some of the people above, the Broadcastify guy for example, and they are, as mentioned, ready doing their own thing so didn't really care about what I wanted to share.

reply

robotastic 12 hours ago [–]

This sounds awesome - If you have any documentation up on how to do this, I would love to point to it from the trunk-recorder wiki.

reply

runawaybottle 1 day ago [–]

I'd love to read a write up on this if you ever feel the urge.

reply

bertmuthalaly 1 day ago [–]

Have you taken a look at the output?

I selected the NYC scanner and found many examples like this:

June 8, 2020 9:03 PM EDT: "Google Launcher new job and I want to play better third-party colder or does the people from the vegetable okay"

reply

artemisyna 1 day ago [–]

Not sure how easy/difficult this would be from an implementation perspective, but perhaps transcription could be provided as a pay-as-you-go service. Ie, users could pay into a pool for live transcription of a stream for a certain period of time or to retroactively transcribe specific streams.

I'd have to imagine that stream listening follows some sort of power (or otherwise 80/20) law, so hopefully that would help with the expense.

reply

coffeeboy27 21 hours ago [–]

Hey! I want to say that I love Broadcastify for listening to my local police and fire scanners so thank you for providing that free service without requiring app downloads and so on. Awesome work!

I will say though I recently noticed that during the recent unrest in my city that the stream had a caption along the lines of 'police are using encrypted channels' which is understandable but disappointing from the perspective of a citizen looking for transparency.

I realize both the possible legal and technical difficulty of implementing something like this but have you had any conversations about how to maybe combat this? Times of unrest are not only a massive opportunity for Broadcastify to try to grow its user base but it's also when transparency is at its peak importance.

reply

optimumsuave 1 day ago [–]

Thank you for the kind words. My team and I simply want to help people and create accountability/transparency.

We decided to ask for forgiveness on this project - we do in fact have a premium account with you guys! ffmpeg records and segments your streams, sox removes silence from scanners, Google transcribes, redis serves it all up quickly!

A broadcastify.com premium account is not a great excuse, but I would love to have a conversation at length with you and your team! How do we get in contact? Thanks again!

reply

robotastic 1 day ago [–]

Feel free to add something like this to some of the feeds from http://openmhz.com I can give you some pointers on how to pull down the streams of audio. They are already chopped into M4A files.

reply

robbie480 1 day ago [–]

Hey Lindsay, I'm the one who just added EBRCS to Calls. Wondering if a solution to both cost problems would be to (optionally) have submitters upload a transcript along with each call? Could build a model into trunk-recorder maybe?

reply

blantonl 1 day ago [–]

That might be an option, but we'd have to somehow get training models down to the client, which is one of the issues.

We're working on an client ingest models now that work on more of a "tasking" perspective, where someone deploys a device that is GPS enabled and then we send an ingest task to fill in coverage, start new coverage, etc. But this is predicated on low cost ingest devices (read: RPi and RTL sticks) which might not have the horsepower needed for transcription at the client level.

reply

jshann8384 22 hours ago [–]

That would be cool, kind of like GroundStation or SatNOGS

reply

godzillabrennus 1 day ago [–]

You could conceivably do this using the text to speech recognition on a Mac using Loopback to capture the stream to a microphone input.

https://rogueamoeba.com/l...

Someone clever enough could create containers to run the software locally and have many loops running off many streams to many instances of the text to text feature.

reply

lunixbochs 1 day ago [–]

Conventional speech recognition will not have very good accuracy from this sort of task out of the box. You will basically need to do local (as in, based on map location) language modelling, and should probably do custom acoustic modeling as well (training a neural network on what radio speech sounds like).

reply

crazygringo 1 day ago [–]

Did you look at the transcripts? They do not have very good accuracy. A better solution probably should do those things; this one I'm assuming is not.

reply

lunixbochs 1 day ago [–]

https://news.ycombinator.c...

> Please don't comment on whether someone read an article. "Did you even read the article? It mentions that" can be shortened to "The article mentions that."

Yes, I looked the website, and I also have my website that does the same thing in a very similar manner. My comment was entirely in response to the suggestion Mac speech recognition be used for this. Based on previous experiments I have personally performed, it would be even worse than the website's accuracy. I then pointed out what a good solution might look like (and neither my website nor the linked website do the good solution yet)

reply

jcims 1 day ago [–]

Assuming you tried Amazon Transcribe? How did it fare?

reply

runawaybottle 1 day ago [–]

Accuracy is a little wonky even with real speech to text toolkits like Kaldi (which I'll mention in a sec to a given get started with it).

I've had some decent results with the following:

https://cmusphinx.github.io/

I have research to see if you can hand tag my own samples to see if that offers significant accuracy improvements (Sphinx? I can't say if it can accurately transcribe one voice consistently).

Google and Watson are not too free, and I I believe Watson has a audio length limit (possibly limited by free tier, or possibly limited in general for all tiers).

Cool to see some real world attempts using this stuff.

reply

cranermon 1 day ago [–]

Not exactly related to this specific post, but do you worry about the slow transition to encrypted channels that emergency departments are making?

reply

rachelbythebay 1 day ago [–]

Example audio: https://www.broadcastify.c...

This is close enough to the Seattle feed that you can compare & contrast.

Heard: "clear my first call ocean nova 470" On the site: "charlie my first call"

So, yeah, this still has a long way to go. I considered and discarded this in 2011 because it was pure insanity, and as another comment suggests, it's highly context-sensitive.

"ECR" is El Camino Real. "Vets" is Veterans Blvd.

But...

"Code 99" is the emergency button - for one department... and it means something else entirely for another, just 20 miles apart.

I'd love to have it, but it still seems out of reach.

reply

jcrsv 1 day ago [–]

Your scanner project was a marvel to me at the time.

Until a system can train right down to the officer/dispatcher I don't see how this gets beyond 50-60% accurate. I can't even transcribe half the calls, not sure why I should expect a computer to (at this time anyway).

reply

3KCalhoun 1 day ago [–]

Like to see a way to play the original audio clip the OCR came from.

Sorry, talking about the project this thread is linked to.

reply

slicuah 1 day ago [–]

Seems like there could be a lot of use cases where there you don't need anything even close to perfection.

reply

ddddddddlost 1 day ago [–]

I really like this idea - it also has the side effect of making all of these streams easily searchable later on.

That being said, transcription still isn't at the point where it can handle radio garble.

> because I'm going to take that group that someone over the Willing. Can you raise Lieutenant lady and having semi-open Buddhist month

> you know the carrot for you and Hudson no accident

Reminds me of the early days of Google Voice voicemail transcription.

reply

tiberenberg 1 day ago [–]

Absolutely agree. I feel like the next step here would be to:

* Put the audio clip next to the transcribed text

* Flag transcribed text which fails some heuristics

* Allows users to suggest edits to the transcription

* Train the transcription models on contributed edits + police specific lingo

reply

gfxxcs 1 day ago [–]

Cooperative correcting editing, genius-style, would be perfect. Big events would lead to incredibly accurate transcripts.

reply

wPsarks 1 day ago [–]

"strength tenets we do not have any to 9 available readable five-year-old advisor. We have several dogs around miracles"

reply

lunixbochs 1 day ago [–]

I built something similar to this, open-source, here: https://feeds.talonvoice.com

It's not currently streaming any feeds because Google Speech is pretty expensive, but I have the expertise and plan to train my own speech models that would be less expensive to run and more accurate than Google at this task as well. (Use my `feeds` site has UI for users to listen to the audio and fix/vote on transcriptions, and correct corrections are propagated quickly to other users.

reply

optimumsuave 1 day ago [–]

Developer of murph.live here - after reading this thread I have some ideas I'd like to vet with everyone here.

1. We need to post links to the source of the stream. I neglected to do that for fear of cease and desist, but now I realize we need to create accountability on our platform. I will be contacting broadcastify.com to ensure we can direct users to a source.

2. We need a disclaimer on the site directly in your face. I agree with everyone here - this could potentially spread misinformation and do more harm than intended. These transcripts should be read with caution. Additional messaging from us is a must.

3. We need a better acoustic model. Google is too much $$$ and although I'm an engineer, I'm not a linguistics machine learning expert. Can anyone help me with this please?

Our mission was to create transparency into our government - not cause harm. There is a lot of responsibility creating a tool like this and we want to get it right.

With that being said, this site blew up in a few hours. I'm overwhelmed. Please let me know if you'd like to help. Thank you to everyone for the feedback so far - it all helps immensely.

reply

robotastic 1 day ago [–]

It might be worth checking out https://www.assemblyai.com they let you build a custom audio model. One challenge with the audio from these radios is that it goes through some heavy compression. Traditional models will have a lot of challenges. Give a system that uses analog audio a try. The quality of the audio is a lot better.

reply

somegudave 1 day ago [–]

I doubt that the google speech model was meant to deal with AMBE compressed voice. I think you will need to create your own speech-to-text model to solve the quality issues.

reply

nellv 1 day ago [–]

When I was studying photojournalism on the side, for a while I ran a vintage radio scanner over the various local police/fire/etc. public simulcast frequencies. It didn't pay off with any "breaking news" leads, but one of the surprise benefits was the strong impression of professionalism and good nature by many.

I should be clear that I know there's tragic racism and other injustices everywhere, and that problems tend to be worse than is visible to most people, and I think now is the time to finally right some wrongs that we should've a long time ago. Perhaps relevant to that, it was really reassuring to observe signs of goodness in the institutions that many people never have occasion to. Problems need to be solved, but there's hope from multiple directions.

reply

belorn 1 day ago [–]

Automatic recording, transcribing and publish. A recipe to force them to change equipment and start using encryption.

I don't know the law in the US, but here in Sweden any police investigation is kept under strict secrecy until completed. When a crime occur and journalist ask for details the answer is always the same. While investigation is ongoing no details may be given. That would not work if there is a searchable transcription of the communication online.

So an nice technical achievement, but the more successful it is the faster it will be made obsolete in terms of getting information out of police scanners.

reply

optimumsuave 1 day ago [–]

Hello Hacker News!!!

I am the developer of murph.live - I just want to thank all of you for taking the time to check it out and give us excellent feedback. I stumbled upon this post and now have goosebumps.

This started by listening to police scanners throughout the night during recent protests in Seattle, WA. I wanted to help and I immediately put my credit card down for Google's Speech to Text API.

As for the inbound streams, @blantonl is spot on - we use the streams from a premium account on broadcastify.com (thank you for not sending a cease and desist yet!).

A few dockerized ffmpeg processes segment the streamed audio into 30 second wav files. Subsequently, sox removes silence from the audio files as police scanners have quite a bit of downtime between transmissions. The performance is very scalable using docker containers to record and trim.

Currently, we pipe these trimmed wav files to the Google Speech API - as others have mentioned this is $$$. We are receiving donations, but this dependency on Google will eventually need to be eliminated.

I have started looking into possible solutions using NLP and other acoustic models to bring the costs down. Honestly, speech processing is not my forte so I'm kind of shooting in the dark here. I am currently testing pre-trained models for wav2letter++, kaldi, vosk, and maybe deepspeech.

We can all agree the quality of the transcripts is something to be desired and improved upon. Potentially dangerous if transcribed incorrectly, but nonetheless we wanted to launch to give citizens a platform to provide transparency into our government. The idea is what counts right now.

Thanks again and I will be responding to a bunch of comments on here! You all rock!

reply

rasz 1 day ago [–]

https://news.ycombinator.c... first comment might be of interest

reply

petee 1 day ago [–]

Nice but mostly incoherent - *"Google Launcher new job and I want to play better third-party colder or does the people from the vegetable okay"*

reply

solarengineer 1 day ago [–]

Given the current global (and US) social situation of mistrust of the police, and of people running wild with information, it may be wise to shut down this website and fix technical issues first.

Otherwise, you do run the risk of "See the screenshots - the police transcriptions have been changed to become police friendly now!"

reply

silverax93 1 day ago [–]

I was reading a post yesterday about how someone found difficult to gauge the level of violence in the current protests. As media is so polarised that you either get shots of peaceful protests or complete mayhem, depending on the paper's bias.

This got me thinking: What if we get the police radio transcripts, and extract every mention of violence? Then have a central hub of drones that get dispatched automatically to any address of "interest" filming.

It's a bit crazy, most likely illegal. But definitely would be cool to see! Arguably truly unbiased "reporting". If you want to know the situation, just plug in to the feeds and get a decent overview of how it looks.

reply

tapmap 1 day ago [–]

Would be great to link to the real recording, so you can hear it for yourself.

reply

pathhandwaving 14 hours ago [–]

I didn't see where anyone pointed this out but scanners are only one way emergency services communicate and they also know people are listening in. More often than not the important conversations are communicated over the phone. It is also important to note that since these agencies know that civilians listen in it can become a way to intentionally misinform the listeners.

reply

nogbit 17 hours ago [–]

This is so wrong it needs to be taken offline. The Translations aren't even close and that's dangerous.

I seriously looked into doing this myself, but for this reason and costs I bailed.

reply

yellow_pos8 1 day ago [–]

Neat! Wonder what service they are using for the transcription as some of the results are pretty rough probably from both poor quality audio and a non-domain lexicon. Custom speech models could probably go a long way here to clean it up and I would bet there's some public interest in contributing transcriptions for training. Even if the site allowed user submitted corrections. Having this indexed and searchable could be powerful.

reply

LVDOVICVS 1 day ago [–]

Thanks for the encouragement in continuing with my court reporting classes. It's still going to be a good career for the foreseeable future.

reply

larrywright 1 day ago [–]

This actually would solve one of my use cases that broadcastify itself does not. I don't listen to the scanner often, and often when I do it's in response to hearing sirens or seeing something. By then, the call has already gone out so tuning in to the scanner isn't all that helpful. Having transcripts of recent calls available would be quite useful.

reply

stuartaxelowen 1 day ago [–]

Very cool! Is this yours, OP? I built a text classifier platform (https://www.taggit.io/) and I'd love to explore building something on top of this text stream to notify people of important happenings around the protests in real time.

reply

illuminated 1 day ago [–]

Not mine, I've got the URL over IM few minutes before posting it here

reply

stuartaxelowen 1 day ago [–]

Thanks!

reply

mirh_ 1 day ago [–]

Is it legal to listen to these in the states?

I know OFCOM (UK FCC) state it's illegal to listen to certain frequencies without permission, which seems delightfully unenforceable (for legal reasons I most definitely do not possess a wideband software defined radio or a similarly sensitive children's toy radios)

reply

jauer 1 day ago [–]

(not a lawyer, if one shows up, listen to them) Generally speaking, you can listen to anything, as long as it's not for commercial gain. The exception to this is mobile phones, because it was trivial to listen to them using children's toy radios.

There's apparently some uncertainty around handling of encrypted emergency services communications: https://www.rtdna.org/cont...

https://www.fcc.gov/consu...

reply

ryanwatkins 1 day ago [–]

Yes, it is legal in the U.S. federal law allows it - P.R. 91-36 (FCC 93-410), though New York law attempts to make it illegal to do in your car: https://codes.findlaw.com/...

reply

inamberclad 1 day ago [–]

The emergency-service voice transmissions in the clear around me are fire and medical. I am presuming that the busy digital signals in the same VHF band are police.

reply

atarian 1 day ago [–]

Apart from the police and people listening out of interest/curiosity, who else taps into police scanners? It seems like a really narrow and competitive niche to be in for developers.

reply

robbie480 1 day ago [–]

According to Broadcastify (https://www.broadcastify.c...), there's 43k listeners across 7000 streams right now

reply

blantonl 1 day ago [–]

It is ridiculously competitive in the mobile app store market, because there is a lot of revenue out there. All of the apps source their streams from Broadcastify, and they license Broadcastify's API and content to do so. The apps are enormously popular, and the market generates hundreds of thousands of dollars a *month* in revenue.

I deal constantly with nefarious developers who often feel like Broadcastify has an obligation to grant them a license. We had a terrible dust up this weekend with a developer that was upset we wouldn't grant him a license to develop "yet another police scanner app" and he went on a warpath with me personally.

The app store environment is a major pain and exhausting...

reply

ketzo 1 day ago [–]

Seems like y'all are definitely positioned right in the market, though, if everyone's using you! Do you do any marketing aimed specifically at developers of these apps?

reply

edyc 1 day ago [–]

Curious on why that is. Can you provide any insights on why it would generate hundreds of thousands a month?

reply

themodelplumber 1 day ago [–]

Local news services.

reply

tobych 1 day ago [–]

A CNN reporter mentioned he was "listening to the police scanner in one ear" the other day while he was reporting live on a protest.

reply

lobster45 1 day ago [–]

Criminals

reply

kube-system 1 day ago [–]

I hear people say this from time to time, but I don't think I've ever heard anything actually said on a scanner that would be plausibly useful to a criminal.

reply

pugworthy 1 day ago [–]

Some of the text though are pretty weak. I have no idea if this is just garbled transcription, or dreamy cop talk...

"movie up and around pain and dreamy I'm not"

or

"happy person that you can just so may I help you"

reply

lunixbochs 1 day ago [–]

It's the transcription, I think this kind of site should allow you to fall back to listening to the audio. Even with high accuracy, if there's no human curation, there will be pretty awkward mistakes every once in a while in the text.

reply

whoisalan 1 day ago [–]

Very useful mapped version in NYC https://scanmap.mobi/NY/

reply

abrichr 1 day ago [–]

This is great. Is it open source? I would love to know more about how it works. Where do you get the source data? What is required in order to add more cities?

Thanks for sharing!

reply

yasoob 1 day ago [–]

I would assume that its piping the audio output from a radio receiver to some live transcription API. Would also love to find out how it actually works.

reply

starpilot 1 day ago [–]

> got another call from the zombie shelter for 24th Avenue got two males waiting inside

Oh shit.

reply

buserarm 1 day ago [–]

Which channel though? In each of these cities there are several frequencies...

reply

vondur 1 day ago [–]

Our local PD is going to be going to encrypted communications soon I believe.

reply

quickthrower2 1 day ago [–]

Looks like DNN speak.

reply

flaque 1 day ago [–]

I'm so glad someone did this

reply

m1bg 1 day ago [–]

Speaking of which, does anyone know the _frequencies_ for police scanners in Seattle? I have an SDR box, but Googling yields nothing of value.

reply